

Lecture 3

HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

Internal Memory Cache

Ch 4, Ch 5 [Sta10]
Key Characteristics
Locality
Cache
Main Memory

(Sta10 Table 4.1)

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 2

Key Characteristics of Memories / Storage

Location	Performance
Processor	Access time
Internal (main)	Cycle time
External (secondary)	Transfer rate
Capacity	Physical Type
Word size	Semiconductor
Number of words	Magnetic
Unit of Transfer	Optical
Word	Magneto-Optical
Block	Physical Characteristics
Access Method	Volatile/nonvolatile
Sequential	Erasable/nonerasable
Direct	Organization
Random	
Associative	

(Sta10 Table 4.1)

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 2

Goals

- I want my memory lightning fast
- I want my memory to be gigantic in size
- Register access viewpoint
 - data access as fast as HW register
 - data size as large as memory
- Memory access viewpoint
 - data access as fast as memory
 - data size as large as disk

cache
HW solution
virtual memory
HW help for SW solution

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 3

Memory Hierarchy

- Most often needed data kept close
- Access to small data sets can be made fast
 - simpler circuits
 - smaller gate delays
- Faster ~ more expensive
- Large can be bigger and cheaper (per byte)

up: smaller, faster, more expensive, more frequent access
down: bigger, slower, less expensive, less frequent access

(Sta10 Fig 4.1)

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 4

Principle of locality (paikallisuus)

- In any given time period, memory references occur only to a small subset of the whole address space
- = The reason why memory hierarchies work

Prob (small data set) = 95% "Cost" (small data set) = 0.01 μs
Prob (the rest) = 5% "Cost" (the rest) = 0.1 μs

Aver cost = 95% * 0.01 μs + 5% * 0.1 μs = 0.015 μs

- Average cost is close to the cost of small data set
- How to determine data for that small set?
- How to keep track of it?

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 5

Principle of locality

- In any given time period
 - Memory references occur only to a small subset of the whole address space
- Temporal locality (ajallinen) ○ ○
 - It is likely that a data item referenced a short time ago will be referenced again soon
- Spatial locality (alueellinen) ○ ○
 - It is likely that a data items close to the one referenced a short time ago will be referenced soon

MEM: 345 23 71 8 305 63 91 2

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 6

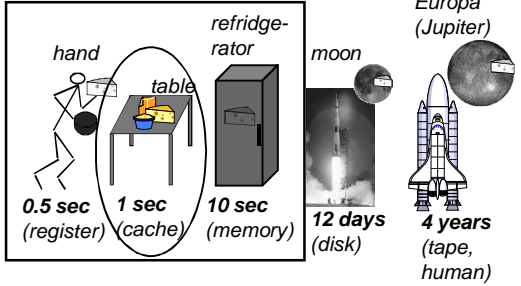
Computer Organization II

Cache

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 7

Teemu's Cheesecake

Register, on-chip cache, memory, disk, and tape speeds relative to times locating cheese for the cheese cake you are baking...



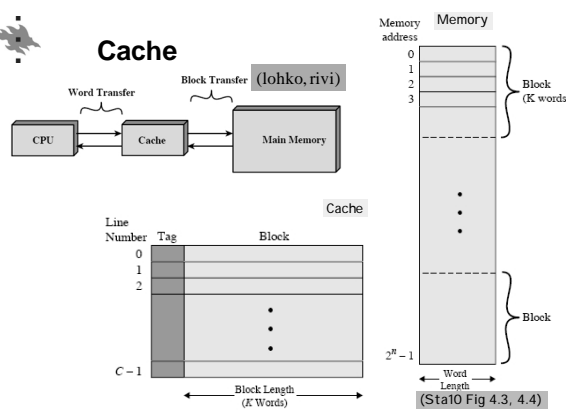
Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 8

Cache Memory (välimuisti)

- How to access main memory as fast as registers?
- Locality → Use (CPU) cache!
 - Keep most probably referenced data in fast cache close to processor, and rest in memory
 - Most of data accesses only to cache
 - hit ratio 0.9-0.99
 - Cache is much smaller than main memory
 - Cache is (much) more expensive (per byte) than memory
- Note: file cache is another thing...

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 9

Cache



Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 10

Cache Read

(RA = Real Address, Main memory address)

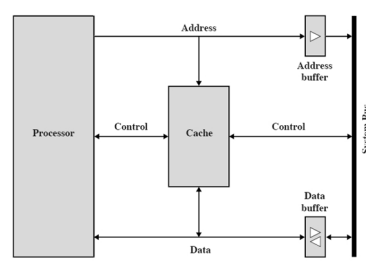
```

    graph TD
        START([START]) --> Receive[Receive address RA from CPU]
        Receive --> IsInCache{Is block containing RA in cache?}
        IsInCache -- "No  
Miss" --> AccessMain[Access main memory for block containing RA]
        IsInCache -- "Yes  
Hit" --> Fetch[Fetch RA word and deliver to CPU]
        Fetch --> DONE([DONE])
        AccessMain --> Allocate[Allocate cache line for main memory block]
        Allocate --> WriteDirty{Write "dirty" cache line back to memory?}
        WriteDirty --> LoadMain[Load main memory block into cache line]
        WriteDirty --> Deliver[Deliver RA word to CPU]
        LoadMain --> Fetch
        Deliver --> Fetch
    
```

(Sta10 Fig 4.5)

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 11

Cache Organization



(Sta10 Fig 4.6)

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 12

Cache Design

Cache Size	Write Policy
Mapping Function	Write through
Direct	Write back
Associative	Write once
Set Associative	Line Size
Replacement Algorithm	Number of caches
Least recently used (LRU)	Single or two level
First in first out (FIFO)	Unified or split
Least frequently used (LFU)	
Random	

(Sta10 Table 4.1)

- Cache Size & Line Size
 - Many blocks help for temporal locality
 - Large blocks help for spatial locality
 - Larger cache is slower
 - Multi-level cache

Typical sizes:
 L1: 8 KB - 64 KB
 L2: 256 KB - 8 MB
 L3: 2 MB - 48 MB

Discussion?

Mapping

- Which block (if any) contains the referenced memory location?
 - Is the block in cache?
 - Where in the cache is it located?
- Solutions
 - Direct mapping (*suora kuvaus*)
 - One possible location
 - Fully associative mapping (*täysin assosiatiivinen*)
 - Any possible location
 - Set associative mapping (*joukkoassosiatiivinen*)
 - Some possible locations

Cache simulation tools:
<http://www.ecs.umass.edu/ece/koren/architecture/Cache/frame0.htm>

Direct Mapping

- Each block has only one possible location (line) in cache
 - determined by index bits
- Several blocks may map into same cache line
 - identified with different tag bits

Block number (in memory): 0x2480, 0x6480, 0xA480

34 bit address (byte address): tag (21), index (8), byte offset (5)

Cache line size = Block size = $2^5 = 32\text{ B}$

Unique bits that are different for each block, Stored into cache line

Fixed location in cache → fixed cache size = $2^8 = 256\text{ blocks} = 8\text{ KB}$

Direct Mapping Example

Word = 4B (here)
 Block size = $2^3 = 8\text{ bytes} = 64\text{ bits}$
 Cache line size

ReadW I2, 0xA4

8 bit address (byte address): tag (2), index (3), offset (3)

tag	index	offset	block, 64b
000:			
001:			
010:			
011:	01		54 A7 00 91 23 66 32 11
100:	11		77 55 55 66 66 22 44 22
101:	01		65 43 21 98 76 65 43 32
110:	10		
111:			

No match

compare

Read new memory block from memory address 0xA0=10100000 to cache location 100, update tag, and then continue with data access

Direct Mapping Example 2

ReadW I2, 0xB4

tag	index	offset	block, 64
000:			
001:			
010:			
011:	01		54 A7 00 91 23 66 32 11
100:	11		77 55 55 66 66 22 44 22
101:	01		65 43 21 98 76 65 43 32
110:	10		00 11 22 33 44 55 66 77
111:			

start with 4th byte

compare → Match

Fully Associative Mapping (6)

- Each block can be in any cache line
 - tag must be complete block number

Alpha AXP uses 34 bit memory addresses

34 bit address (byte address): tag (29), offset (5)

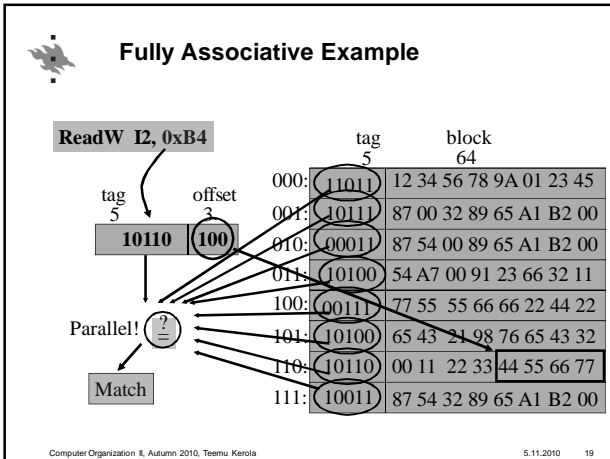
Block number (in memory)

Offset from the beginning of the block (in bytes)

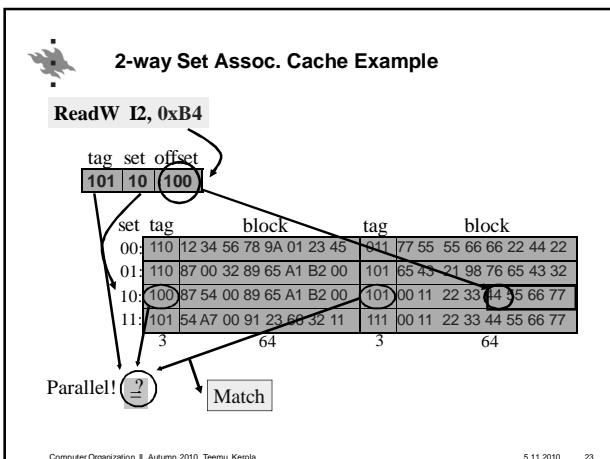
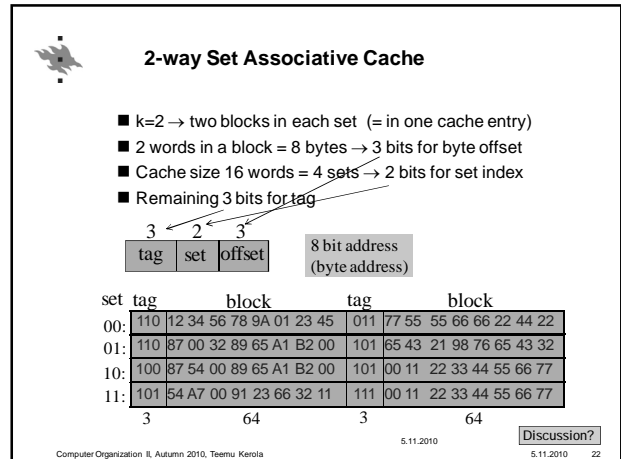
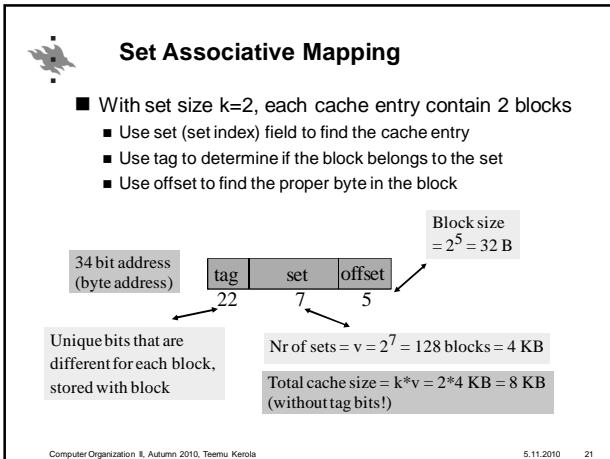
Block size = $2^5 = 32\text{ B}$

Unique bits that are different for each block

Each block can be anywhere Cache size can be any number of blocks



- ### Fully Associative Mapping
- Lots of circuits
 - tag fields are long - wasted space?
 - each cache line tag must be compared in parallel with the memory address tag
 - lots of wires, comparison circuits
 - large surface area on chip
 - Final comparison "or" has large gate delay
 - did any of these 64 comparisons match?
 - $\log_2(64) = 6$ levels of binary OR-gates
 - how about 262144 comparisons?
 - 18 levels?
 - \Rightarrow Can use it only for small caches



- ### Set Associative Mapping
- Set associative cache with set size $k=2$ = 2-way cache (common)
 - Degree of associativity = nr of blocks in a set = v
 - Large degree of associativity?
 - More data items in one set
 - Less "collisions" within set
 - Final comparison (matching tags?) gate delay?
 - Maximum (nr of cache lines)
 - \Rightarrow fully associative mapping Whole cache is one set!
 - Minimum (1)
 - \Rightarrow direct mapping Each cache line is a set!

Cache Replacement Algorithm

- Which cache block to replace to make room for new block from memory?
- Direct mapping: trivial
- First-In-First-Out (FIFO)?
- Least-Frequently-Used (LFU)?
- Random?
- Which one is best / possible?
 - Chip area?
 - Fast? Easy to implement?

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 25

Cache Write Policy – memory writes?

- Write through (*läpikirjoittava*)
 - Each write goes always to cache and memory
 - Each write is a cache miss!
- Write back (*lopuksi/takaisin kirjoittava*)
 - Each write goes only to cache
 - Write cache block back to memory
 - only when it is replaced in cache
 - Memory may have stale (old) data
 - cache coherence problem (*eheys, yhdenmukaisuus*)
- Write once (*"vain kerran kirjoittava?"*)
 - Write-invalidate Snoopy-cache coherence protocol for multiprocessors
 - Write invalidates data in other caches
 - Write to memory at replacement time, or when some other cache needs it (has read/write miss)

Coherence problems:
More users of the same data: memory valid? cache valid? multiple processors with own caches

A bit set

Discussion?

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 26

Cache Line Size

- How big cache line?
- Optimise for temporal or spatial locality?
 - bigger cache line → better for spatial locality
 - more cache lines → better for temporal locality
- Best size varies with program or program phase?
- Best size different with code and data?
- 2-8 words?
 - word = 1 float??

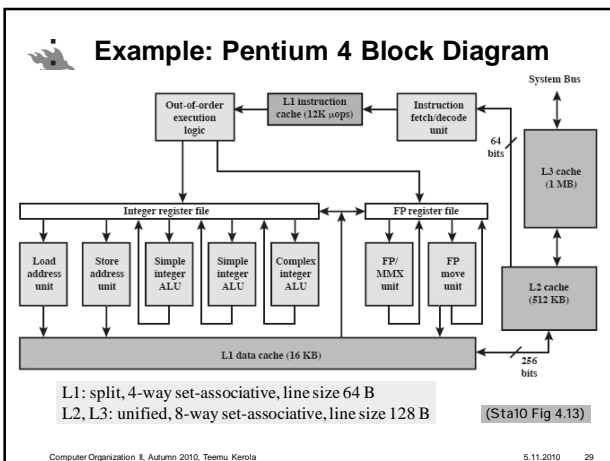
Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 27

Types and Number of Caches

- Same cache for data and code, or not?
 - Data references and code references behave differently
 - Unified vs. split cache (*yhdistetty/erilliset*)
 - Split cache: can optimise structure separately for data and code
- One cache too large for best results?
 - "Smaller is faster"
- Multiple levels of caches
 - L1 on same chip as CPU
 - L2 on same package or chip as CPU
 - older systems: same board
 - L3 on same board as CPU

Trend towards split caches: Pentium, Power PC, ARM..

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 28



Computer Organization II

Main Memory

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 30

Main Memory Types (katoava, haihtuva)

Memory Type	Category	Erasure	Write Mechanism	Volatility
Random-access memory (RAM)	Read-write memory	Electrically, byte-level	Electrically	Volatile
Read-only memory (ROM)	Read-only memory	Not possible	Masks	Nonvolatile
Programmable ROM (PROM)			Electrically	
Erasable PROM (EPROM)	UV light, chip-level			
Electrically Erasable PROM (EEPROM)	Read-mostly memory	Electrically, byte-level		
Flash memory		Electrically, block-level		

(Sta10 Table 5.1)

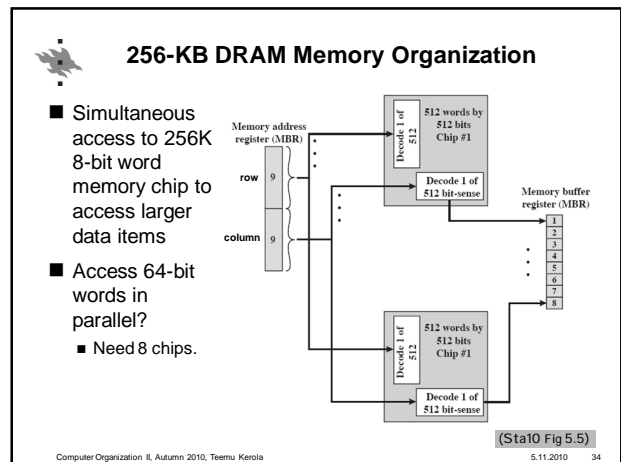
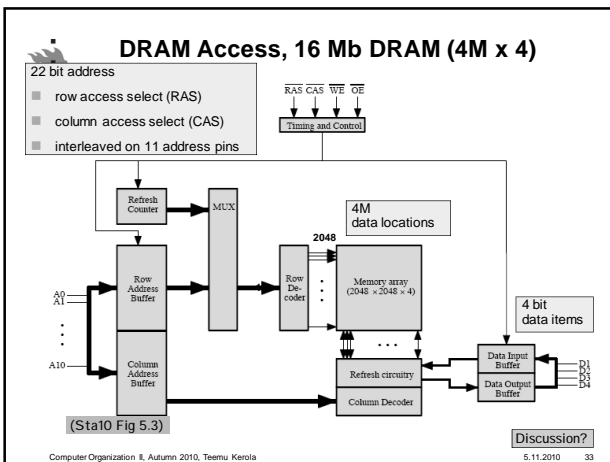
- Random access semiconductor memory
 - Direct access to each memory cell
 - Access time same for all cells

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 31

RAM

- Dynamic RAM, DRAM
 - Periodic refreshing required Analog: Charge on capacitors
 - Refresh required after read
 - Simpler, slower, denser, bigger (bytes per chip)
 - Access time ~ 60 ns (?)
 - Main memory? (early systems)
- Static RAM, SRAM
 - No periodic refreshing needed Digital: flip-flop gates
 - Data remains until power is lost
 - More complex (more chip area/byte), faster, smaller
 - Access time ~ 25 ns (?)
 - Cache?

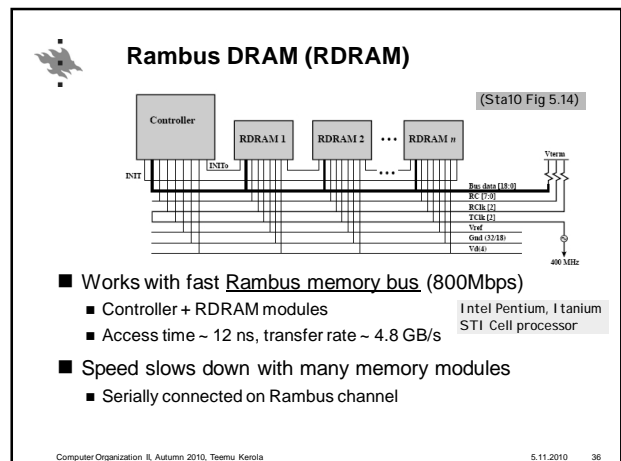
Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 32



SDRAM (Synchronous DRAM)


- CPU clock synchronizes also the bus
 - Runs on higher clock speeds than ordinary DRAM
 - CPU knows how long it takes to make a reference, can do other work while waiting
- 16 bits in parallel
 - Access 4 DRAMs (4 bits each) in parallel
 - Access time ~ 18 ns, transfer rate ~ 1.3 GB/s
- DDR SDRAM, double data rate
 - Current main memory technology
 - Supports transfers both on rising and falling edge of the clock cycle
 - Consumes less power
 - Access time ~ 12 ns, transfer rate ~ 3.2 GB/s

Computer Organization II, Autumn 2010, Teemu Kerola 5.11.2010 35



MRAM

- Magnetoresistive Random Access Memory (MRAM)
 - Data stored with magnetic fields on two plates
 - Magnetic field directions determine bit value
- Non-volatile, data remains with power off
 - Fast to read/write
 - No upper limit for write counts (Flash has upper limit)
 - Access time comparable to DRAM, almost as fast as SRAM
- Future open
 - Small market share now
 - Expensive now (2006: \$25 4Mbit, 2008: \$15 4Mbit, Freescale)
 - Still under development
 - May replace flash "in a few years"
 - Has not by 2010 ...



Computer Organization II, Autumn 2010, Teemu Kerola

5.11.2010 37

Summary

- Memory hierarchy
- Cache
 - Size, Line size, Mapping, nr of levels, unified or split, write policy, replacement policy
- Memory
 - DRAM, SRAM
 - Overall features, no details required
 - Current technologies
 - Synchronous DRAM (SDRAM)
 - Double Data Rate DRAM (DDR)
 - Rambus DRAM

Computer Organization II, Autumn 2010, Teemu Kerola

5.11.2010 38/38

Review questions

- Memory hierarchy and principle of locality?
- Different ways to use locality in cache solutions?
- Differences of associative and set associative mappings?
- Why to have separate caches for instructions and data ?

Computer Organization II, Autumn 2010, Teemu Kerola

5.11.2010 39